

Two-Stream Gated Fusion ConvNets for Action Recognition

Jiagang Zhu^{*†}, Wei Zou^{*‡}, Zheng Zhu^{*†}

^{*}Institute of Automation, Chinese Academy of Sciences, Beijing, China

[†]University of Chinese Academy of Sciences, Beijing, China

[‡]TianJin Intelligent Tech.Institute of CASIA Co., Ltd

Email: zhujiagang2015@ia.ac.cn, wei.zou@ia.ac.cn, zhuzheng2014@ia.ac.cn

Abstract—The two-stream ConvNets in action recognition always fuse the two streams' predictions by the weighted averaging scheme. This fusion way with fixed weights lacks of pertinence to different action videos and always needs trial and error on the validation set. In order to enhance the adaptability of two-stream ConvNets, an end-to-end trainable gated fusion method, namely gating ConvNet, is proposed in this paper based on the MoE (Mixture of Experts) theory. The gating ConvNet takes the combination of convolutional layers of the spatial and temporal nets as input and outputs two fusion weights. To reduce the over-fitting of gating ConvNet caused by the redundancy of parameters, a new multi-task learning method is designed, which jointly learns the gating fusion weights for the two streams and learns the gating ConvNet for action classification. With the proposed gated fusion method and multi-task learning approach, competitive performance is achieved on the video action dataset UCF101.

I. INTRODUCTION

Human action recognition is important for applications of human-robot interaction, behavior analysis and surveillance. Early works [1], [2], [20] utilized hand-crafted spatial-temporal local descriptors and powerful encoding methods. Inspired by the successes of deep learning for image classification [18], lots of works have explored deep convolutional neural networks (CNN) [3], [4] for video classification and achieved higher performance than hand-crafted methods recently.

This paper mainly focuses on improving the performance of the two-stream ConvNets [3], [4], [8], [9], [16], [25] in action recognition. The two-stream ConvNets [3] contain the spatial net and the temporal net, which take RGB frames and consecutive optical flow stacks as inputs respectively. The predictions of the two streams are always fused by evenly averaging or weighted averaging [3], [4]. This fixed weight fusion method cannot make the best use of the capacity of the spatial and the temporal nets. Because each of them fires on different aspects of videos: the spatial net focuses on the appearance and scene contents of videos, while the temporal one on the motion. Also, different video frames of the same and the different classes contain different amount of spatial and temporal cues. Fusing the predictions of the spatial and the temporal nets with fixed weight may not capture the contents of videos well and always needs trial and error on the validation set. Some fusion methods [8], [9] have also been proposed for the two-stream ConvNets, but they are not

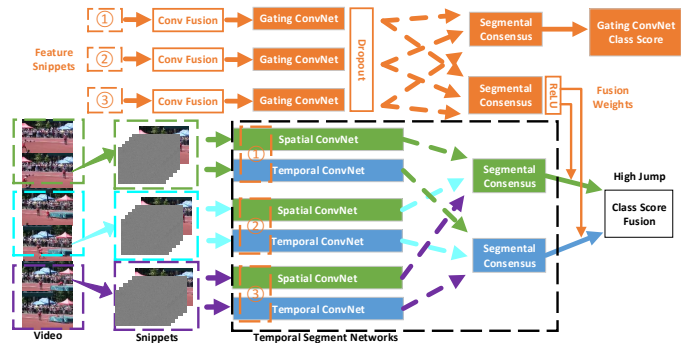


Fig. 1. Gated TSN: Newly added gating ConvNet are in orange color. Its inputs are called feature snippets, where number ①, ②, ③ denote the different segment level feature maps from the two streams. Each of feature snippets is the combination of feature maps from the same layer of the spatial and the temporal nets. There are two independent fully connected layers for the gating ConvNet, one for learning the gating fusion weights for the two streams, the other for the action classification of the gating ConvNet.

designed for fusing the predictions between the two streams. The SCI (Sparsity Concentration Index) fusion [8] gives a weighted score scheme according to the sparsity degrees of the crop-level prediction. The SCI fusion in their work is used to fuse the predictions of different crops from a single spatio-temporal stream, while our work is for fusing the predictions of the two streams. The conv fusion [9] also fuses the two streams in feature level to get a spatio-temporal stream and a temporal stream. But they fuse the predictions of the spatio-temporal and temporal stream by evenly averaging.

Motivated by above observations, in this paper, an end-to-end trainable gated fusion method is proposed to obtain the adaptive fusion weights for the spatial and the temporal nets in prediction level. We gain insights from the Mixture of Experts (MoE) [10], which are usually consisted of a gating network and more than two expert networks. Each expert network is gated via a Softmax function of the gating network [15] [13]. By gating on a number of experts, they aim to adaptively combine a subset of networks to make prediction. However, there are some differences between the above MoE architectures and our work. Firstly, different from these MoE methods with more than two experts, our work aims to best utilize the spatial and the temporal nets of the two-stream based action recognition methods. Secondly, the inputs for

different experts in previous MoE methods are usually from the same source or a small subset of the same source. While our gating network takes the combination of convolutional feature maps of the spatial and the temporal nets as inputs. It is termed as **gating ConvNet** for it is mainly composed of several convolutional layers. Our fusion method based on the gating ConvNet is termed as **gated fusion**. Besides, a new multi-task learning [3], [12] method is proposed, which jointly learns the adaptive fusion weights for the two streams and the gating ConvNet for action classification. Different from the weighted averaging fusion with fixed weights, the gated fusion is a sample specific fusion method because the gating ConvNet makes a reasonable assignment in two fusion weights for the spatial and the temporal nets adaptively according to the properties of different video inputs.

The main contributions of this paper can be summarized as follows: 1) An end-to-end trainable gated fusion method is proposed for the two-stream ConvNets. 2) A new multi-task learning method is designed, which jointly learns the gating fusion weights for the two streams and the gating ConvNet for action classification. With this approach, the performance of our MoE is improved. 3) With our gated fusion method and multi-task learning approach, competitive performance is achieved on the video action dataset UCF101.

II. APPROACH

In this section, the framework of the gated fusion method will be introduced. Then three aspects of learning the gating ConvNet are detailed.

A. Gated Fusion for the two-stream ConvNets

The spatial and the temporal nets of TSN are selected to act as two experts in our basic framework for their simple architectures and good performance in the two-stream ConvNets. The reader can refer to [4] for more details of TSN. Fig. 1 shows our **gated TSN**. It contains TSN and gating ConvNet. The gating ConvNet takes the combination of convolutional feature maps of the two streams as inputs. Besides, it has two independent fully connected layers, one for learning the gating fusion weights for the two streams and the other for action classification. So they can work in a multi-task learning manner. Through segmental consensus [4] of the gating ConvNet, the video-level fusion weights and the video-level gating ConvNet predictions for action recognition are obtained. The gating ConvNet outputs the gating fusion weights as follows

$$\mathbf{G}_g = \mathcal{H}_g(\mathcal{G}_g(\mathcal{F}_g(f_1; \mathbf{W}_g), \mathcal{F}_g(f_2; \mathbf{W}_g), \dots, \mathcal{F}_g(f_K; \mathbf{W}_g))) \quad (1)$$

where f_k , $k = 1, \dots, K$ is the feature snippets generated by the combination of feature maps of the spatial and the temporal nets. K is number of segments. $\mathcal{F}_g(f_k; \mathbf{W}_g)$ is the function representing the gating ConvNet with parameters \mathbf{W}_g which operates on the feature snippet f_k . \mathcal{G}_g aggregates the frame level fusion weights to get the video-level fusion

weights. Average pooling is adopted for \mathcal{G}_g . \mathcal{H}_g is a ReLU function ensuring the non-negativity of video-level fusion weights. With the gated fusion method, adaptive weighted function (2) of the two streams is obtained, where w_1 and w_2 are the fusion weights outputted by the gating ConvNet $w_1 = G_{g1}, w_2 = G_{g2}$. It is also worth noting that the predictions of the two streams \mathbf{G}_{rgb} and \mathbf{G}_{flow} are fused before Softmax normalization [4], [23] in our gated fusion method

$$\mathbf{G}_{adap} = w_1 \mathbf{G}_{rgb} + w_2 \mathbf{G}_{flow} \quad (2)$$

where \mathbf{G}_{adap} is the weighted prediction.

For the classification branch of the gating ConvNet, the prediction function is defined as

$$\mathbf{G}_c = \mathcal{G}_c(\mathcal{F}_c(f_1; \mathbf{W}_c), \mathcal{F}_c(f_2; \mathbf{W}_c), \dots, \mathcal{F}_c(f_K; \mathbf{W}_c)) \quad (3)$$

where $\mathcal{F}_c(f_k; \mathbf{W}_c)$ is the function representing the gating ConvNet with parameters \mathbf{W}_c which operates on the feature snippet f_k . Note that \mathbf{W}_c and \mathbf{W}_g share parameters except for the fully connected layers. \mathcal{G}_c aggregates frame level predictions into the video-level predictions. Average pooling is adopted for \mathcal{G}_c . The final loss function for the gating ConvNet is

$$\begin{aligned} \mathbf{L} &= \mathcal{L}(y, \mathbf{G}_{adap}) + \lambda \mathcal{L}(y, \mathbf{G}_c) \\ &= - \sum_{i=1}^C y_i \left(G_{adap_i} - \log \sum_{j=1}^C \exp G_{adap_j} \right) \\ &\quad - \lambda \sum_{i=1}^C y_i \left(G_{c_i} - \log \sum_{j=1}^C \exp G_{c_j} \right) \end{aligned} \quad (4)$$

where C is the number of action classes and y_i the ground truth label of class i . $\mathcal{L}(y, \mathbf{G}_{adap})$ is the loss with respect to the predictions of the two streams after gated fusion. $\mathcal{L}(y, \mathbf{G}_c)$ is the loss of the classification branch of the gating ConvNet. λ is the loss weight for classification loss of the gating ConvNet. Standard cross-entropy loss is employed for these two losses respectively. In the back-propagation process, the gradients of the gating ConvNet parameters \mathbf{W} with respect to the loss can be derived as

$$\begin{aligned} \frac{\partial \mathbf{L}}{\partial \mathbf{W}} &= \frac{\partial \mathcal{L}(y, \mathbf{G}_{adap})}{\partial \mathbf{W}_g} + \lambda \frac{\partial \mathcal{L}(y, \mathbf{G}_c)}{\partial \mathbf{W}_c} = \frac{\partial \mathcal{L}}{\partial \mathbf{G}_{adap}} \\ &\quad \left(\mathbf{G}_{rgb} \frac{\partial \mathbf{G}_{adap}}{\partial G_{g1}} \frac{\partial G_{g1}}{\partial \mathcal{H}_g} \frac{\partial \mathcal{H}_g}{\partial \mathcal{G}_g} \sum_{k=1}^K \frac{\partial \mathcal{G}_g}{\partial \mathcal{F}_g(f_k)} \frac{\partial \mathcal{F}_g(f_k)}{\partial W_g} \right. \\ &\quad \left. + \mathbf{G}_{flow} \frac{\partial \mathbf{G}_{adap}}{\partial G_{g2}} \frac{\partial G_{g2}}{\partial \mathcal{H}_g} \frac{\partial \mathcal{H}_g}{\partial \mathcal{G}_g} \sum_{k=1}^K \frac{\partial \mathcal{G}_g}{\partial \mathcal{F}_g(f_k)} \frac{\partial \mathcal{F}_g(f_k)}{\partial W_g} \right) \\ &\quad + \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{G}_c} \sum_{k=1}^K \frac{\partial \mathcal{G}_c}{\partial \mathcal{F}_c(f_k)} \frac{\partial \mathcal{F}_c(f_k)}{\partial \mathbf{W}_c} \end{aligned} \quad (5)$$

B. Implementations of learning the gating ConvNet

Output activation function for the gating ConvNet. In our MoE method with $N = 2$ experts (spatial and temporal net), the outputs of gating ConvNet perform as the confidence ratio between the two streams. To model gating outputs g as a function of its inputs \mathbf{x} , different functions could be considered. ReLU

$$g(x_i) = \max(0, x_i), i = 1, 2 \quad (6)$$

is selected as the output activation function because of its non-negativity and fast convergence speed [18]. It can perform the same role as Softmax [10]

$$g(x_i) = \frac{\exp(x_i)}{\sum_{i=1}^N \exp(x_i)}, i = 1, 2 \quad (7)$$

ReLU could make sense as long as the two output fusion weights $g(x_i), i = 1, 2$ of the gating ConvNet do not become zero together for a specific sample.

Inputs for the gated TSN. The inputs for the spatial and the temporal nets of the gated TSN follow the original TSN, where RGB frames and optical flow stacks are randomly sampled from each of K segments of a video. Each group of RGB frames and optical flow stack starts from the same point in a video. The gating ConvNet in this paper takes the feature maps of different layers of two streams with **concatenation fusion** or **conv fusion** [9] as inputs. Concatenation fusion $\mathbf{y}^{cat} = f^{cat}(\mathbf{x}^a, \mathbf{x}^b)$ stacks the two feature maps $\mathbf{x}^a, \mathbf{x}^b$ from the two streams at the same spatial locations i, j across the feature channels d

$$y_{i,j,2d}^{cat} = x_{i,j,d}^a, y_{i,j,2d-1}^{cat} = x_{i,j,d}^b \quad (8)$$

where $\mathbf{y}^{cat} \in \mathbb{R}^{H \times W \times 2D}$. Conv fusion $\mathbf{y}^{conv} = f^{conv}(\mathbf{x}^a, \mathbf{x}^b)$ first stacks the two feature maps $\mathbf{x}^a, \mathbf{x}^b$ at the same spatial locations i, j across the feature channels d as above equation (8) and subsequently convolves the stacked data with a bank of filters $\mathbf{f}^{1 \times 1} \in \mathbb{R}^{1 \times 1 \times 2D \times D}$ and biases $\mathbf{b} \in \mathbb{R}^D$

$$\mathbf{y}^{conv} = \mathbf{y}^{cat} * \mathbf{f}^{1 \times 1} + \mathbf{b} \quad (9)$$

where $\mathbf{y}^{conv} \in \mathbb{R}^{H \times W \times D}$. The filter $\mathbf{f}^{1 \times 1}$ is used to reduce the dimensionality of concatenation feature maps by a factor of two and is able to model weighted combinations of the two feature maps from two streams at the same spatial location.

Multi-Task Learning for the gating ConvNet. The gating ConvNet takes the convolutional layers of BN-Inception [7] after the above mentioned input fusion layer as its feature extractor, followed by a dropout layer and two independent fully connected layers. Note that the same feature extractor of action classification is used for the gating ConvNet. However, the dimension of classification output (101 for UCF101) is much larger than the dimension of the gating fusion weights (in the case of two-stream, 2). Thus, the gating ConvNet is equipped with redundant degrees of freedom in its feature extractor. Learning this task could be cumbersome [17] and

suffers from severe risk of over-fitting. To relieve over-fitting, a action classification branch is added on top of the final convolutional layer of the gating ConvNet and it could behave as a regularizer for the task of learning the gating fusion weights. At this point, the gating ConvNet has two independent fully connected layers, one for the gating fusion weights, the other for the action classification. These two fully connected layers share the same input layer (inception5b). It is expected that joint learning of the fusion weights and classification could improve the accuracy of our MoE.

III. EXPERIMENTS

A. Dataset and Implementation Details

Experiments are conducted on the standard action dataset: UCF101 [11]. The UCF101 dataset contains 101 action classes and 13,320 video clips. Three training/testing splits are used for evaluation. All experiments are implemented with Caffe [6] and one NVIDIA GTX TITAN X GPU is used for training and testing. Codes will be available at <https://github.com/zhujiagang/gating-ConvNet-code>.

Network Training. The TSN in our gated TSN adopts training strategies including cross modality pre-training, partial BN, dropout and data augmentation [4]. The number of the snippets K is set to 3 for both TSN and gating ConvNet. The loss weight λ is set to 0 when we only learn gating fusion weights, and is set to 1 when we jointly learn gating fusion weights and the gating ConvNet for classification. For the gating ConvNet, the ImageNet pre-training is used and the dropout ratio of dropout out layer is set to 0.8. The mini-batch SGD algorithm is used to learn the network parameters. The training procedures of the gated TSN mainly include three stages: 1) Firstly the two streams of the TSN are trained; 2) Then parameters of these two streams are fixed and we only fine-tune the gating ConvNet for learning the gating fusion weights; 3) When there is no more increase in accuracy, lastly we do joint learning of the gating fusion weights and gating ConvNet for action classification. Training the gating ConvNet consumes much more memory than training the spatial and temporal nets respectively, so a smaller batch size is needed (set to 4) than the first training stage (set to 32). L2 norm of gradients is clipped at 40 and momentum term is set to 0.9. For training the gating ConvNet, the learning rate is initialized as 0.001 and decreases to 0.0001 when there is no more increase in accuracy. The model is selected by early stopping. Optical flows are extracted by the TVL1 optical flow algorithm [24]. RGB frames and optical flows are extracted from videos in advance.

Network Testing. For each video during testing, 25 RGB frames and optical flow stacks are sampled. Meanwhile, the crops of 4 corner and 1 center, and their horizontal flippings are obtained from sampled frames. Each pair of RGB frame and optical flow stack starts from the same point in a video. For each pair of them, the gated fusion for the spatial and temporal net is applied. All predictions of crops in a video are averaged to get a video-level result.

TABLE I
ACCURACY (%) OF THE GATED TSN WITH SOFTMAX AND RELU AS THE OUTPUT ACTIVATION FUNCTION OF THE GATING CONVNET RESPECTIVELY ON THE UCF101 (SPLIT 1).

Input for gating ConvNet	Softmax	ReLU
Conv fusion of inception4e	94.05	94.11
Concatenation fusion of inception4e	93.97	93.96

TABLE II
ACCURACY (%) OF THE GATED TSN WHEN DIFFERENT INPUT LAYERS WITH CONCATENATION OR CONV FUSION ARE USED FOR THE INPUTS OF THE GATING CONVNET ON THE UCF101 (SPLIT 1).

Input layer for gating ConvNet	concatenation fusion	conv fusion
conv1	93.97	93.96
conv2	93.91	93.54
inception3c	93.93	93.89
inception4e	93.96	94.11
inception5b	93.89	93.87

B. Ablation Studies

ReLU or Softmax. For the output activation function of the gating ConvNet, two activation functions are explored: Softmax and ReLU. Feature maps with concatenation fusion or conv fusion [9] of inception4e from two streams are taken as the inputs of the gating ConvNet. As shown in Table I, the gating ConvNet with ReLU is comparable to that with Softmax in the accuracy of the gated TSN. It can be seen that ReLU could perform the same role as Softmax when used as the final gating activation function. It is also found that the gating ConvNet with ReLU as the final activation function converges faster than that with Softmax. This complies with the fact that ReLU could have faster training speed than saturating nonlinearities when used as the layer activation function [18]. In later experiments, ReLU is adopted for the gating output activation function.

Different input layers and ways of fusing these layers. Each pair of convolutional layers of the two streams (conv1, conv2, inception3c, inception4e, inception5b) with concatenation or conv fusion is chosen for each experiment. The performance of different input layers for the gating ConvNet is summarized in Table II. While different fusion methods (conv, concatenation) and different input layers perform slightly different, the inception4e as the input layer of the gating ConvNet with conv fusion gets the highest accuracy. This is different from the previous work about two-stream feature fusion [9], where the fusion of the highest convolutional layers after ReLU gains the best result.

Different network architectures. To test the generality of our gated fusion for the two-stream ConvNets in action recognition, we also do experiments on different network architectures including CaffeNet [6] and VGG16 [5]. Two

TABLE III
COMPARISON OF THE ACCURACY (%) AMONG DIFFERENT FUSION METHODS FOR DIFFERENT NETWORK ARCHITECTURES AND DIFFERENT TWO-STREAM METHODS ON THE UCF101 (SPLIT 1).

Architectures	gated fusion	weighted ave	SCI
CaffeNet (two-stream)	71.80	71.75	71.34
VGG16 (two-stream)	78.94	78.82	77.61
CaffeNet (TSN)	74.93	74.55	70.86
VGG16 (TSN)	88.02	87.95	85.00
BN-Inception (TSN)	94.11	93.81	93.96

kinds of two-stream based methods, namely the original two-stream ConvNets [3] and TSN [4] are implemented with these architectures. The number of the snippets K during training is set to 3 for TSN and 1 for the original two-stream ConvNets. Different fusion methods such as weighted averaging fusion with fixed weight, fusion based on SCI (Sparsity Concentration Index) [8] and our gated fusion method are used for all these networks. Results are summarized in Table III. For weighted averaging fusion, the predictions of the two streams before Softmax normalization are fused and the best weight is selected with grid search on the validation set for each experiment. As the method in [8] with SCI fusion has only one stream, in our two-stream method with SCI fusion, average fusion is added in stream level after its crop level probability fusion. For the spatial and the temporal nets with CaffeNet and VGG16, it is found that these two networks suffer from severe over-fitting in UCF101 due to the limited training data. The gated fusion can not do better than the weighted averaging fusion in already over-fitting expert networks [13]. To reduce the over-fitting of the CaffeNet and VGG16 in the two-stream ConvNets, all their fully connected layers are removed. Then the training schemes of the original two-stream and TSN are followed to get the final spatial and temporal models. As shown in Table III, the gated fusion always performs the best in different architectures and different two-stream methods, which shows the advantage of assigning the gating fusion weights for the two streams. The SCI fusion performs comparably well with the weighted averaging fusion.

Multi-task learning for the gating ConvNet. Further, a different fully connected layer for action classification is added on top of the last convolutional layer of the gating ConvNet. The network is fine-tuned on the previous trained gating ConvNet by jointly learning the gating fusion weights and the action classification. As shown in Table IV, after adding a classification branch, the accuracy of the gated TSN increases by 0.08%, 0.36% and 0.7% on the three splits of UCF101 respectively. So, it can be concluded that learning the gating fusion weights could benefit from learning the gating ConvNet for action classification.

C. Comparison with the State of the Art

After above analysis of the gating ConvNet, final experiments on all three splits of UCF101 are implemented with our proposed methods. Mean average accuracy on three test

TABLE IV
ACCURACY (%) OF THE GATED TSN WHEN THE GATING CONVNET DOES
JOINT LEARNING OF ACTION CLASSIFICATION AND LEARNING THE
GATING FUSION WEIGHTS ON THE UCF101 (THREE SPLITS).

UCF101 split	gated fusion	+gating ConvNet classification
split 1	94.11	94.19
split 2	94.12	94.48
split 3	94.14	94.84

TABLE V
COMPARISON OF THE ACCURACY (%) OF OUR GATED TSN WITH OTHER
STATE-OF-THE-ART METHODS.

Methods	UCF101
IDT	85.9
MoFAP	88.3
Two-stream ConvNet	88.0
C3D (3 nets)	85.2
FstCN	88.1
LTC	91.7
ST-ResNet	93.4
TSN(2 modalities)	94.0
TSN(3 modalities)	94.2
Gated TSN(ours)	94.5
DOVF	94.9
TLE	95.6

sets of UCF101 is calculated as the final result. As shown in Table V, the gated TSN are compared with both traditional approaches [2], [20] and deep learning methods [3], [4], [21], [22], [25]–[27]. It is noted that our gated TSN only employs 2 modalities (RGB frame and optical flow stacks) as inputs and improves upon the original TSN with 2 modalities by 0.5%. It even exceeds the TSN with 3 modalities by 0.3%. This improvement demonstrates that the weighted averaging fusion with fixed weight could not fully exploit the capacity of the two streams on different samples, even with three streams, while the TSN with our gated fusion method could improve performance by adaptively assigning the fusion weights to different streams.

D. Network Visualization

In Fig. 2, the distributions of the fusion weights for the gated TSN with and without our multi-task learning are displayed, corresponding histograms of the fusion weights for the spatial net are followed. None of the coordinate of the points on the first two subplots is zero, implying that the gating ConvNet has learnt that combining the spatial and the temporal nets is better than that only with single network or no network. With our multi-task learning, the output points of the gating ConvNet distribute more sparsely than that without multi-task learning. It could also be observed in the last subplot that the fusion weights for the spatial net range from 0.4 to 0.7. It is more wider than that without multi-task learning, whose fusion weights for the spatial net are mostly centered between

0.5 and 0.65. This may account for the 0.36% increase of the accuracy on UCF101 split2 in Table IV. With our joint learning method, an adaptive selection space is expanded for assigning the fusion weights for the two streams with more variations according to the current inputs.

Finally, some examples of the classification results of the gated TSN and the weighted averaging fusion are shown and compared in Fig. 3. In the first three subplots of Fig. 3, the spatial stream always has high confidence for the ground truth label, while the temporal stream has high confidence for the incorrect class. In these cases, the higher fusion weights for the temporal stream than the spatial one may weaken the confidence to the ground truth, may even lead to prediction failures just as shown in these three examples. Our gated fusion assigns the spatial stream higher weights than the temporal one in all these three cases, which gives correct predictions with higher confidence, and proves it has learned that the spatial stream should be trusted more in these cases. It is also noticed that in the fourth subplot of Fig. 3, the ground truth label, namely MoppingFloor, is not predicted into the top-5 by both the two streams, but after both fusion methods, it appears again. Fusing the predictions of the two streams with our gated fusion brings the result to be true by giving higher weights to the spatial stream than the temporal one.

IV. CONCLUSION

In this work, an end-to-end trainable gated fusion method is proposed for the two-stream ConvNets in action recognition. Besides, it is shown that our joint learning of the gating fusion weights for the two streams and learning the gating ConvNet for action classification is helpful in improving the accuracy of the gated TSN. Our techniques in this work could also be extended to the semantic segmentation domain, where multi-stream deep neural networks are employed [17].

ACKNOWLEDGMENT

This work is supported in part by the National High Technology Research and Development Program of China under Grant No.2015AA042307, the National Natural Science Foundation of China under Grant No.61773374, and in part by Project of Development In Tianjin for Scientific Research Institutes Supported By Tianjin government under Grant No.16PTYJGX00050.

REFERENCES

- [1] I. Laptev, "On space-time interest points," *IJCV*, vol. 64, no. 2, pp. 107–123, 2005.
- [2] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *ICCV*, 2013.
- [3] K. Simonyan, A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *NIPS*, 2014.
- [4] L. Wang *et al.*, "Temporal segment networks: towards good practices for deep action recognition," in *ECCV*, 2016.
- [5] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [6] Y. Jia *et al.*, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.
- [7] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015.




Fig. 2. The distribution of the fusion weights for the gated TSN (Top row), where horizontal axis and vertical axis represent the ReLU output of gating ConvNet for the temporal net and the spatial net respectively. Corresponding histograms of the fusion weights (Bottom row) for the spatial net where horizontal axis represents the gating fusion weights for the spatial net and vertical axis stands for the number of samples. All on the validation set of UCF101 (split 2).

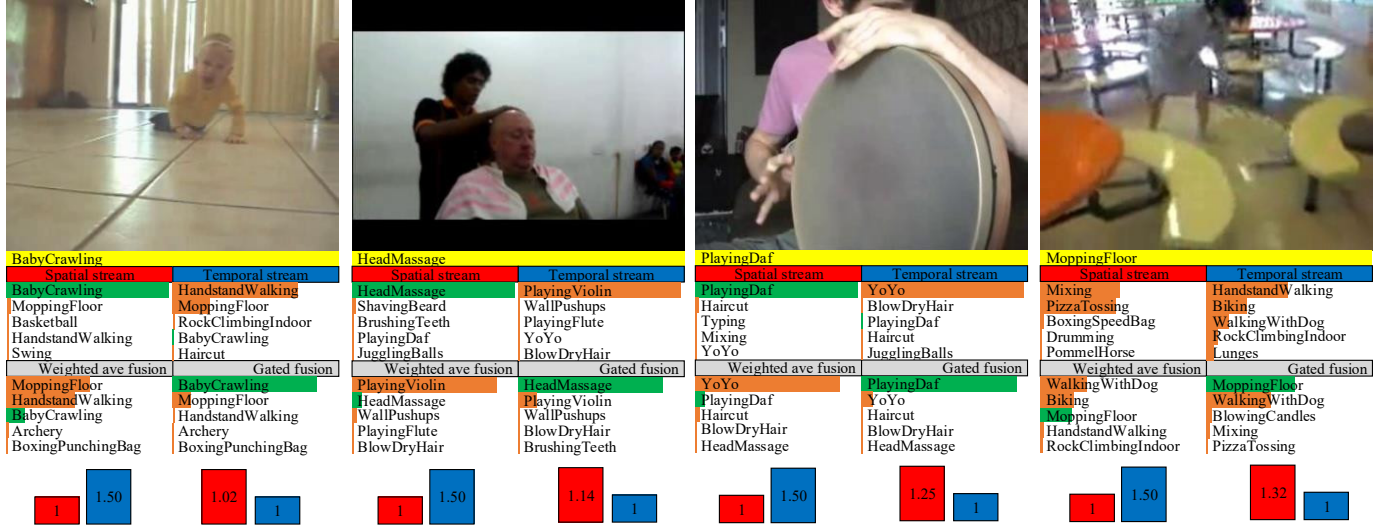


Fig. 3. A comparison of top-5 predictions between the weighted averaging fusion with fixed weights and our gated fusion method on UCF101. The yellow bars stand for the ground truth label. The predictions of single stream are also shown, red for the spatial stream and blue for the temporal stream. Green and orange bars indicate correct and incorrect predictions respectively and the length of each bar shows its confidence. The predictions after different fusion methods and their fusion weights (confidence ratio) in the bottom of the figure by different fusion method are also displayed correspondingly.

- [8] L. Sun, K. Jia, D. Y. Yeung, and B. E. Shi, "Human Action Recognition using Factorized spatio-Temporal Convolutional Networks," in *ICCV*, 2015.
- [9] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional Two-Stream Network Fusion for Video Action Recognition," in *CVPR*, 2016.
- [10] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79-87, 1991.
- [11] K. Soomro, A. R. Zamir, and M. Shah, "UCF101: A dataset of 101 human actions classes from videos in the wild," *arXiv preprint arXiv:1212.0402*, 2012.
- [12] S. Ren, K. He, R. Girshick, and J. Sun., "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *NIPS*, 2015.
- [13] V. Ramamurti and J. Ghosh, "On the use of Localized Gating in Mixture of Experts Networks," in *Applications and Science of Computational Intelligence*, 1998.
- [14] L. V. D. Maaten and G.E. Hinton, "Visualizing high-dimensional data using t-SNE," *JMLR*, vol. 9, pp. 2579-2605, 2008.
- [15] N. Shazeer *et al.*, "Outrageously large neural networks: the sparsely-gated mixture-of-experts layer," in *ICLR*, 2017.
- [16] R. Girdhar *et al.*, "ActionVLAD: Learning spatio-temporal aggregation for action classification," in *CVPR*, 2017.
- [17] A. Valada, J. Vertens, A. Dhall, and W. Burgard, "AdapNet: Adaptive Semantic Segmentation in Adverse Environmental Conditions," in *ICRA*, 2017.
- [18] A. Krizhevsky, I. Sutskever, and G.E. Hinton, "ImageNet classification with deep convolutional neural networks," in *NIPS*, 2012.
- [19] R. Girshick, "Fast R-CNN," in *ICCV*, 2015.
- [20] L. Wang, Y. Qiao, and X. Tang, "MoFAP: A multi-level representation for action recognition," *IJCV*, vol. 119, no. 3, pp. 254-271, 2016.
- [21] G. Varol, I. Laptev, and C. Schmid, "Long-term temporal convolutions for action recognition," *arXiv preprint arXiv:1604.04494*, 2016.
- [22] D. Tran *et al.*, "Learning spatiotemporal features with 3D convolutional networks," in *ICCV*, 2015.
- [23] L. Wang, Y. Xiong, D. Lin, and L. V. Gool, "UntrimmedNets for Weakly Supervised Action Recognition and Detection," in *CVPR*, 2017.
- [24] C. Zach, T. Pock, and H. Bischof, "A duality based approach for realtime tv-L1 optical flow," in *29th DAGM Symposium on Pattern Recognition*, 2007.
- [25] A. Diba, V. Sharma, and L. V. Gool, "Deep temporal linear encoding networks," in *CVPR*, 2016.
- [26] Z. Lan, Y. Zhu, and A. G. Hauptmann, "Deep Local Video Feature for Action Recognition," *CVPR workshop*, 2017.
- [27] C. Feichtenhofer, A. Pinz and R. P. Wildes, "Spatiotemporal Residual Networks for Video Action Recognition," in *CVPR*, 2017.
- [28] T.-Y. Lin, A. RoyChowdhury, and S. Maji, "Bilinear cnn models for fine-grained visual recognition," in *ICCV*, 2015.
- [29] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell, "Compact bilinear pooling," in *CVPR*, 2016.